Docket No.:

IN THE U.S. PATENT AND TRADEMARK OFFICE

In re patent application of:
P.S. Chaporkar

Serial No.: (Goes here)

Filed: (Date of Filing)

For: (Title goes here)

Group Art Unit: (Goes Here)

Examiner: (Goes here)

Assistant Commissioner for Patents
Washington, D.C. 20231

SUBMISSION OF PRIORITY DOCUMENT

Sir:

Submitted herewith is a certified copy of India Application Number 727/Del/02 filed on 9th July 2002, upon which application the claim for priority is based.

Respectfully submitted,

John M Kelly

John M. Kelly
Reg. No. 33,920

The Law Firm of John Kelly
2117 L Street, NW #347
Washington, DC 20037
Phone (202) 744-4787

सत्यमेव जयते

INTELLECTUAL PROPERTY **INDIA**

GOVERNMENT OF INDIA
MINISTRY OF COMMERCE & INDUSTRY,
PATENT OFFICE, DELHI BRANCH,
W - 5, WEST PATEL NAGAR,
NEW DELHI - 110 008.

I, the undersigned, being an officer duly authorized in accordance with the provision of the Patent Act, 1970 hereby certify that annexed hereto is the true copy of the Application, Provisional and Complete Specification and Drawing Sheets filed in connection with Application for Patent No. ***727/Del/02*** *dated 9th July* ***2002****.*

Witness my hand this 28th Day of May 2003.

(S.K. PANGASA)
Assistant Controller of Patents & Designs

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

FORM 1

0727-2

09 JUL 2002

THE PATENTS ACT, 1970
(39 OF 1970)
APPLICATION FOR GRANT OF PATENT
(See section 5(2), 7, 54 and 135 and rule 33A)

Govt. of India Patent Office
New Delhi
Received Rs. 5000/- in cash.
Cheque/M.O./T.P.O./D.D.
on 9/7/02
Vide Entry No. 2661 in the Register of Valuables
O. ....Cashier

1. We, **HIMACHAL FUTURISTIC COMMUNICATIONS LIMITED**, a joint stock company incorporated under the Indian Companies Act, 1956, of 8, COMMERCIAL COMPLEX, MASJID MOTH, GREATER KAILASH, NEW DELHI 110 048.

2. Hereby declare:-

a) that we are in possession of an invention titled **"A METHOD FOR FAST COST-EFFECTIVE INTERNET NETWORK TOPLOGY DESIGN".**

b) that the Provisional Specification relating to this invention is filed with this application.

c) that there is no lawful ground of objection to the grant of patent to us.

ORIGINAL

3. Further declare that the inventors for the said invention are:-

(a) ANURAG KUMAR, Professor, ECE Department, Indian Institute of Science, Bangalore - 560012, Indian National.

(b) PRASANNA S. CHAPORKAR, Project Associate, Centre for Electronics Design and Technology, Indian Institute of Science, Bangalore - 560012, Indian National.

(c) NANDITA D., Project Associate, Centre for Electronics Design and Technology, Indian Institute of Science, Bangalore - 560012, Indian National.

(d) JOY KURI, Assistant Professor, Centre for Electronics Design and Technology, Indian Institute of Science, Bangalore - 560012, Indian National.

4. We claim the priority from the application filed in convention countries, particulars of which are as follows:-

nil nil nil

5. We state that the said invention is an improvement in or modification of the invention, the particulars of which are as follows and of which I am the applicant / patentee:

nil nil nil

6. We state that the application is divided out of our application, the particulars of which are given below and pray that this application deemed to have been filed on_____(nil) under Section 16 of the Act.

nil nil nil

7. That we are the assignee or legal representative of the true and first inventors.

8. That our address for service in India is as follows:-

AGARWAL JETLEY & CO.

A-2/78, SAFDARJUNG ENCLAVE,

NEW DELHI 110029

9. Following declaration was given by the inventors or applicants in the convention country:

nil nil nil

10. That to the best of our knowledge, information and belief the fact and matters stated herein are correct and that there is no lawful ground of objection to the grant of patent to us on this application.

11. Following are the attachment with the application:

Forms 1,2,3 and 62

We request that a patent may be granted to us for the said invention.

Dated this 8th day of July 2002

SHAILENDRA SINGH, Advocate,
AGARWAL JETLEY & CO.
A-2/78, SAFDARJUNG ENCLAVE,
NEW DELHI 110029
AGENTS FOR THE APPLICANTS

To
The Controller of Patents,
The Patent Office,
at New Delhi.

FORM 2
THE PATENTS ACT, 1970
(39 OF 1970).

PROVISIONAL SPECIFICATION.
(See Section 10)

0727-2

09 JUL 2002

"A METHOD FOR FAST COST-EFFECTIVE INTERNET NETWORK TOPLOGY DESIGN "

HIMACHAL FUTURISTIC COMMUNICATIONS LIMITED.

HIMACHAL FUTURISTIC COMMUNICATIONS LIMITED, a joint stock company incorporated under the Indian Companies Act, 1956, of 8, COMMERCIAL COMPLEX, MASJID MOTH, GREATER KAILASH, NEW DELHI - 110048.

The following specification particularly describes the nature of the invention and the manner in which it is to be performed

ORIGINAL

The invention relates to a method for fast cost-effective internet network toplogy Design.

A Method for Fast, Cost-Effective Network Topology Design.

Network design and management are two fundamental problems faced by Internet Service Providers (ISPs). For a new ISP, the network design problem consists of the following:

- Deciding where to place the routers: the Points of Presence (PoPs) of the ISP
- Interconnecting the routers: the topology of the network
- Determining the capacities of the interconnecting links

- Determining aggregate flow routing: the path to be taken by aggregate flows between source routers and destination routers

For an ISP trying to design a network, the natural granularity for measuring traffic is aggregate flows and not individual session micro-flows. Thus, routing refers to routing of aggregates between source-destination pairs.

Deciding where the PoPs should be is mostly an administrative issue, where resource and infrastructure availability are primary concerns. For example, PoPs are mostly located in metropolitan cities where resource and infrastructure levels are higher than in towns and villages. Once the PoP locations have been decided, a new ISP needs predictions of the traffic that will be flowing between pairs of PoPs. This document does not consider the problems of PoP location and traffic prediction; these are assumed to be given.

The **design objective** of a new ISP is to obtain a network that costs the least while supporting the predicted traffic requirements. The network cost is determined by (a) *cost of the routing equipment* (routers, switches and associated equipment), (b) *cost of the interconnecting links* which are normally leased from the bandwith provider (for example, BSNL). In this document, the cost of the routing equipment is treated as fixed, and the focus is on reducing the cost of the interconnecting links. The link costs are governed by two factors: (i) the distance of a link joining the routers, (ii) the capacity of the link.

Management of an existing ISP network may also call for capacity expansion planning. In this scenario the ISP already has a network deployed, but the growth of traffic necessitates capacity expansion. In this case, the ISP is faced with the questions: (a) whether to add a link between two routers, and if so, what capacity link to choose, and (b) whether to upgrade an existing link to one of higher capacity. The design objective in this case is to upgrade the network at least cost.

The network design problem is often more than merely obtaining a *feasible* network at least cost. Additional constraints arise from the following requirements:

1. *Guaranteeable performance*: Network design may be influenced by some performance guarantees that the ISP may want to announce as part of its service offering (for example: a minimum average assured rate for file

downloads through its network). This may require the use of higher capacity links.

2. *No flow splitting*: An aggregate flow from a source router to a destination router must be carried along a single path only, and cannot be split among multiple paths. This constraint may be imposed by routing technology (for example, OSPF weights).

3. *Symmetric routing*: The aggregate flow from router *i* to router *j* and the aggregate flow from router *j* to router *i* must use the same path in opposite directions.

4. *Hop-constrained routing*: The maximum number of hops on the path of any aggregate flow must not be greater than a specified *hopcount*.

5. *Survivability*: Occasional failures of routers or links should not cause interruptions in the service provided. This indicates the need to provide spare capacity in the network; for example, there should be at least 2 independent paths between any pair of routers.

The problem of finding the minimum cost network, given the cost and traffic requirement matrices, can be formulated as an integer linear program. This problem has been shown to be NP-complete. Since a direct solution is computationally very expensive, efforts have been made to obtain "indirect" but useful information about the problem; for example, upper and lower bounds to the optimal cost have been obtained in some cases. Heuristic and approximate solutions have also been tried.

The *uncapacitated* network design problem, which is concerned with providing connectivity among the nodes at minimum cost, has been addressed in [1]. Each link has a certain cost associated with it; the cost depends on the distance covered by the link. Since this is an uncapacitated problem, traffic matrices are not given. A tight lower bound to the optimal cost is obtained by a dual ascent procedure.

The problem of capacitated network design for multicommodity flows is considered in [2] and [3]. Capacities are available in multiples of a base unit. A lower bound to the optimal cost is given by a branch and bound heuristic. [2] also gives an upper bound based on polyhedral cutting plane techniques.

Several other papers consider the same problem. [4] considers ε -approximation techniques for this problem. When capacities are restricted to 2 values only, [5] develops an iterative algorithm based on facet defining inequalities as cutting planes. Optimization techniques for designing survivable networks are discussed in [6,7,8, 9,10].

It can be seen that the works above address a variety of problems; however, each is a simplified version of the general topology design problem of practical interest. The general problem can include some or all of the additional requirements listed in points (1) to (5) above. Even the simplified versions are hard problems to solve; so throwing in additional requirements makes the problem extremely hard. Designing heuristic algorithms is the only practical option.

One such heuristic, named MENTOR (MEsh Network Topology Optimization and Routing), is proposed in [11]. It is a heuristic to find a minimum cost network that supports the given traffic matrix. However, there are several limitations: (a) link capacities are restricted to 1 value only, (b) the cost of a link obeys the triangle inequality, (c) flow splitting to an arbitrary granularity is assumed. Additional requirements like survivability and hop constraints cannot be incorporated. Also, the algorithm cannot start with any partial topology; so problems like cost effective capacity expansion cannot be addressed. This motivates the need for a more flexible heuristic.

Proposed Solution: A heuristic algorithm to design network topologies satisfying given traffic requirements, hop and connectivity constraints *at low and possibly least cost* is proposed.

The information available to the algorithm is as follows:

1. PoPs: this gives the positions of the nodes and the distance between any 2 of them.

2. Traffic matrix: this gives the predicted amount of traffic between any 2 nodes.

3. A set of all the available link capacities.

4. Cost matrix: this is a 3-dimensional matrix that gives the cost of a link of capacity C between any 2 nodes; the value C is chosen from the set of available capacities.

Given n nodes, there can be at most $n(n-1)$ directed links in the network. These can be thought of as "potential" links, in the sense that the algorithm chooses some links from this set of potential links for interconnecting the given nodes. Each potential link has 2 variables associated with it: *current_capacity* and *current_traffic*. For a fresh network design problem, *current_capacity* and *current_traffic* are initialised to zero for each potential link. For a network upgradation problem, each link in the present topology has its *current_capacity* and *current_traffic* initialised to its present value, while for the rest of the potential links, each variable is initialised to zero.

The algorithm can now be described as follows:

1. Obtain an ordered sequence of (source, destination) pairs.

2. Choose the first (source, destination) pair from the sequence. For each potential link, consider the minimum capacity that would be required *if* the predicted traffic for that (source, destination) pair was made to flow through that link. This minimum capacity is chosen from the set of available capacities. Define the *differential cost* of the link as the cost difference between the required minimum capacity and the present capacity of the link.

3. Using the Bellman-Ford algorithm, find the least-cost hop-constrained path from the source to the destination with the link metrics being the differential costs. This path will be used to route the traffic for the (source, destination) pair.

4. For every link on the selected path, update the *current_capacity* and the *current_cost*. If the (source, destination) pair was the last in the sequence, then stop. Else, omit this (source, destination) pair from the sequence and go to Step 2.

At any iteration, the aim of the algorithm is to augment the current topology such that the predicted traffic requirement between the chosen (source, destination) pair can be supported with minimum additional cost.

No specific criteria for ordering the node pairs is given in Step 1. Several orderings are possible and each such ordering leads to a variant of the algorithm. Usually, the cost of links is seen to be concave with respect to both distance and the capacity of the link. A chosen ordering should seek to utilize this concavity.

The algorithm described above is a *greedy* algorithm in the sense that it seeks to minimize the incremental cost at each iteration. However, it may be possible to further improve the cost of the topology. Two methods to do this are given below.

The first method, named "Link Removal Heuristic" (LRH), attempts to reduce the overall cost by further refining the topology. The basic idea is to explore possible reductions in cost by removing a link at a time. The aggregate flows passing through the link to be removed need to be re-routed through the remaining links. If such removal of a link and the consequent re-routing result in decreased topology cost, then the link is removed. These steps are repeated for all the links.

The second method, named "Flow Removal Heuristic" (FRH), attempts to reduce cost in a different way. Suppose the original topology cost was Γ_{old}. The idea is to remove a flow from the network and obtain a possibly modified topology with possibly less cost. Now, with this modified topology as a starting topology, we attempt to put back and route the flow that was removed. This may result in a new topology (since additional capacity may be required on some links) with a cost Γ_{new} (say). If Γ_{new} is less than Γ_{old}, then the new topology is retained. All flows are examined one by one in the above way.

The issue of survivable topologies is, in general, quite complicated. We restrict our attention to single link failures only. The requirement is to have alternate paths for connections passing through a failed link. Clearly, this implies that additional capacity must be installed on the links of the network and the objective is to do this at least cost.

The algorithm proposed for obtaining topologies that survive single link failures is based on the algorithm ...

is provided by the output of the basic algorithm. Suppose one link in this topology fails. Then the traffic flowing though that link must be re-routed using other links. The basic algorithm is again used at this point to explore different paths that can carry the traffic passing through the failed link, and the least cost path is chosen. The above process is repeated for every link in the starting topology.

To estimate the quality of the proposed method, several example problems in topology design were considered. For small problems, it is possible to obtain the *optimal* cost by solving an Integer Linear Program. The cost of the design obtained from the method can then be compared with the optimal cost to see how well the method performs. It was seen that, for all the small problems for which the optimal costs were computable, there was no difference between the optimal cost and the cost of the design given by the method. This indicates that the method performs well. A key advantage in using the method is the speed with which designs are obtainable. This is shown in the Figure 1 attached. It can be seen that the time taken to obtain the optimal solution increases drastically as the network size increases (a characteristic of all NP-hard problems), and attempting to find the optimal solution for large network design problems is impractical

Figure 1 illustrates the time taken by the proposed method to obtain a network topology design versus the network size. The time taken to obtain the optimal solution is also shown.

We claim:

1. A fast method for obtaining a low—cost, and possibly least cost, backbone network forming the basic infrastructure of an Internet Service Provider (ISP) --- the said network consisting of the routers and the links administered and operated by the ISP --- over a given set of router locations, given a traffic demand matrix, the said traffic originating from the demands of the ISP's customers, and a link distance—bandwidth cost matrix, the said cost matrix being given by a bandwidth provider, with

(a) the method being based on sequential addition of traffic demands, each demand being routed over a minimum weight path, where all possible links given the router locations are considered and the link weights are assigned, for each demand, by a differential cost analysis,

(b) the method, in addition, ensuring that a demand between any ingress router and egress router (referred to as an "ingress--egress pair") is not split,

(c) the method, in addition, ensuring symmetric routing of demands between any ingress--egress pair,

(d) the method, in addition, ensuring hop-constrained routing, viz., that a demand between any ingress--egress pair traverses at most a pre-specified number of hops.

2. Two methods, named Link Removal Heuristic (LRH) and Demand Removal Heuristic (DRH), that start with the network provided by the "Method 1." described above and reduce cost by further refinements to the network, the said refinements being obtained by the following steps:

a) removal of a link (demand) from the network,

b) updation of the network such that the removed link (demand) can be re-inserted into the network at least incremental cost.

3. A method named Surviving Single Link Failures Heuristic (SSLFH) that starts with the network provided by the "Method 1." described above and incorporates survivability constraints in the design and thereby ensures sufficient spare capacity that is used in the event of single link failures, the said method consisting of a repeated application of the following steps:

a. removal of a link from the network,

b. re-routing of all demands that were routed through the above link such that the re-routing can be accomplished at least incremental cost.

Dated this 8th day of July 2002.

SHAILENDRA SINGH, Advocate,
AGARWAL JETLEY & CO.
A-2/78, SAFDARJUNG ENCLAVE,
NEW DELHI 110029
AGENTS FOR THE APPLICANTS

To
The Controller of Patents,
The Patents Office,

0727 2 TOTAL SHEETS -1

SHEET 1 OF 1

09 JUL 2002

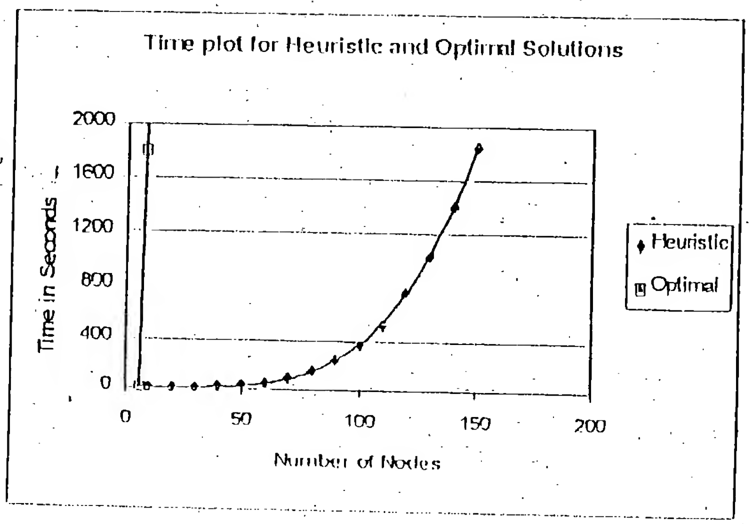

Percentage difference between costs obtained from the heuristic solution and optimal solution is 0 for small number of nodes

Figure 1. The time taken by the proposed method to obtain a network topology design versus the network size

ORIGINAL

SHAILENDRA SINGH, Advocate,
AGARWAL JETLEY & CO.
A-2/78, SAFDARJUNG ENCLAVE,
NEW DELHI 110029
AGENTS FOR THE APPLICANTS

12 MAR 2003

FORM 2

THE PATENTS ACT, 1970
(39 OF 1970)

COMPLETE SPECIFICATION
(See Section 10)

"A METHOD FOR FAST COST-EFFECTIVE INTERNET NETWORK TOPLOGY DESIGN"

HIMACHAL FUTURISTIC COMMUNICATIONS LIMITED

HIMACHAL FUTURISTIC COMMUNICATIONS LIMITED, a joint stock company incorporated under the Indian Companies Act, 1956, of 8, COMMERCIAL COMPLEX, MASJID MOTH, GREATER KAILASH, NEW DELHI-110048

The following specification particularly describes the nature of the invention and the manner in which it is to be performed:-

The invention relates to a method for fast cost-effective internet network topology Design:

A Method for Fast, Cost-Effective Network Topology Design.

Network design and manage are two fundamental problems faced by Internet Service Providers (ISPs). For a new ISP, the network design problem consists of the following:

- Deciding where to place the routers: the Points of Presence (PoPs) of the ISP
- Interconnecting the routers: the topology of the network
- Determining the capacities of the interconnecting links
- Determining aggregate flow routing: the path to be taken by aggregate flows between source routers and destination routers.

For an ISP trying to design a network, the natural granularity for measuring traffic is aggregate flows and not individual session micro-flows. Thus, routing refers to routing of aggregates between source-destination pairs.

Deciding where the PoPs should be is mostly an administrative issue, where resource and infrastructure availability are primary concerns. For example, PoPs are mostly located in metropolitan cities where resource and infrastructure levels are higher than in towns and villages. Once the PoP locations have been decided, a new ISP needs

predictions of the traffic that will be flowing between pairs of PoPs. This document does not consider the problems of PoP location and traffic prediction; these are assumed to be given.

The **design objective** of a new ISP is to obtain a network that costs the least while supporting the predicted traffic requirements. The network cost is determined by (a) cost of the routing equipment (routers, switches and associated equipment), (b) cost of the interconnecting links which are normally leased from the bandwidth provider (for example, BSNL). In this document, the cost of the interconnecting links. The link costs are governed by two factors: (i) the distance of a link joining the routers, (ii) the capacity of the link.

Management of an existing ISP network may also call for capacity expansion planning. In this scenario the ISP already has a network deployed, but the growth of traffic necessitates capacity expansion. In this case, the ISP is faced with the questions: (a) whether to add a link between two routers, and if so, what capacity link to choose, and (b) whether to upgrade an existing link to one of higher capacity. The design objective in this case is to upgrade the network at least cost.

The network design problem is often more than merely obtaining a feasible network at least cost. Additional constraints arise from the following requirements:

1. *Guaranteeable performance*: Network design may be influenced by some performance guarantees that the ISP may want to announce as part of its service offering (for example: a minimum average assured rate for file downloads through its network). This may require the use of higher capacity links.

2. *No flow splitting*: An aggregate flow from a source, router to a destination router must be carried along a single path only, and cannot be split among multiple paths. This constraint may be imposed by routing technology (for example, OSPF weights).

3. *Symmetric routing*: The aggregate flow from router/to router/and the aggregate flow from router/to router/must use the same path in opposite directions.

4. *HOP-constrained routing*: The maximum number of hops on the path of any aggregate flow must not be greater than a specified hopcount.

5. *Survivability*: Occasional failures of routers or links should not cause interruptions in the service provided. This indicates the need to provide spare capacity in the network; for example, there should be at least 2 independent paths between any pair of routers.

The problem of finding the minimum cost network, given the cost and traffic requirement matrices, can be formulated as an integer linear program. This problem has been shown to be NP-complete. Since a direct solution is computationally very expensive, efforts have been made to obtain "indirect" but useful information about the problem; for example, upper and lower bounds to the optimal cost have been obtained in some cases. Heuristic and approximate solutions have also been tried.

The uncapacitated network design problem, which is concerned with providing connectivity among the nodes at minimum cost, has been addressed in [1]. Each link has a certain cost associated with it; the cost depends on the distance covered by the link. Since this is an uncapacitated problem, traffic matrices are not given. A tight lower bound to the optimal cost is obtained by a dual ascent procedure.

The problem of capacitated network design for multicommodity flows is considered in [2] and [3]. Capacities are available in multiples of a base unit. A lower bound to the optimal cost is given by a branch and bound heuristic. [2] also gives an upper bound based on polyhedral cutting plane techniques.

Several other papers consider the same problem. [4] considers approximation techniques for this problem. When capacities are restricted to 2 values only, [5] develops an iterative algorithm based on facet defining inequalities as

cutting planes. Optimization techniques for designing survivable networks are discussed in [6, 7, 8, 9, 10].

It can be seen that the works above address a variety of problems; however, each is a simplified version of the general topology design problem of practical interest. The general problem can include some or all of the additional requirements listed in points (1) to (5) above. Even the simplified versions are hard problems to solve; so throwing in additional requirements makes the problem extremely hard. Designing heuristic algorithms is the only practical option.

One such heuristic, named MENTOR (MEsh Network Topology Optimization and Routing), is proposed in [11]. It is a heuristic to find a minimum cost network that supports the given traffic matrix. However, there are several limitations: (a) link capacities are restricted to 1 value only (b) the cost of a link obeys the triangle inequality, (c) flow splitting to an arbitrary granularity is assumed. Additional requirements like survivability and hop constraints cannot be incorporated. Also, the algorithm cannot start with any partial topology; so problems like cost effective capacity expansion cannot be addressed. This motivates the need for a more flexible heuristic.

Proposed Solution: A method to design network topologies satisfying given traffic requirements, hop and connectivity constraints at low and possibly least cost is proposed.

The information available to the method is as follows:

1. PoPs: this gives the positions of the nodes and the distance between any 2 of them.
2. Traffic matrix: this gives the predicted amount of traffic between any 2 nodes.
3. A set of all the available link capacities.
4. Cost matrix: this is a 3-dimensional matrix that gives the cost of a link of capacity C between any 2 nodes; the value C is chosen from the set of available capacities.

Given n nodes, there can be at most n(n-1) directed links in the network. These can be thought of as "potential" links, in the sense that the algorithm chooses some links from this set of potential links for interconnecting the given nodes. Each potential link has 2 variables associated with it: current_capacity and current_traffic are initialised to zero for each potential link. For a network upgradation problem, each link in the present topology has its current_capacity and current_traffic initialised to its present value, while for the rest of the potential links, each variable is initialised to zero.

The method can now be described as follows:

1. Obtain an ordered sequence of (source, destination) pairs.

2. Choose the first (source, destination) pair from the sequence. For each potential link, consider the minimum capacity that would be required if the predicted traffic for that (source, destination) pair was made to flow through that link. This minimum capacity is chosen from the set of available capacities. Define the differential cost of the link as the cost different between the required minimum capacity and the present capacity of the link.

3. Using the Bellman-Ford method, find the least-cost hop-constrained path from the source to the destination with the link metrics being the differential costs. This path will be used to route the traffic for the (source, destination) pair.

4. For every link on the selected path, update the current_capacity and the current_cost. If the (source, destination) pair was the last in the sequence, then stop. Else, omit this (source, destination) pair from the sequence and go to Step 2.

At any iteration, the aim of the method is to augment the current topology such that the predicted traffic requirement between the chosen (source, destination) pair can be supported with minimum additional cost.

No specific criteria for ordering the node pairs is given in Step 1. Several orderings are possible and each such ordering leads to a variant of the method. Usually, the cost of links is seen to be concave with respect to both distance and the capacity of the link. A chosen ordering should seek to utilize this concavity.

The method described above is a greedy method in the sense that it seeks to minimize the incremental cost at each iteration. However, it may be possible to further improve the cost of the topology. Two methods to do this are given below.

The first method, named "Link Removal Heuristic" (LRH), attempts to reduce the overall cost by further refining the topology. The basic idea is to explore possible reductions in cost by removing a link at a time. The aggregate flows passing through the link to be removed need to be re-routed through the remaining links. If such removal of a link and the consequent re-routing result in decreased topology cost, then the link is removed. These steps are repeated for all the links.

The second method, named "Flow Removal Heuristic" (FRH), attempts to reduce cost in a different way. Suppose the original topology cost was ...T old. The idea is to remove a flow from the network and obtain a possibly modified topology with possibly less cost. Now, with this modified topology as a starting topology, we attempt to put back and route the flow that was removed. This may result in a new topology (since additional capacity may be required on some links) with a cost T_{new} (say). If T_{new} is less than T_{old}, then the new topology is retained. All flows are examined one by one in the above way.

The issue of survivable topologies is, in general, quite complicated. We restrict our attention to single link failures only. The requirement is to have alternative paths for connections passing through a failed link. Clearly, this implies that additional

capacity must be installed on the links of the network and the objective is to do this at least cost.

The method proposed for obtaining topologies that survive single link failures is based on the method described in steps (1) to (4) above. The starting topology is provided by the output of the basic method. Suppose one link in this topology falls. Then the traffic flowing though that link must be re-routed using other links. The basic method is again used at this point to explore different paths that can carry the traffic passing through the failed link, and the least cost path is chosen. The above process is repeated for every link in the starting topology.

To estimate the quality of the proposed method, several example problems in topology design were considered. For small problems, it is possible to obtain the optimal cost by solving an Integer Linear Program. The cost of the design obtained from the method performs. It was seen that, for all the small problems for which the optimal costs were computable, there was no difference between the optimal cost and the cost of the design given by the method. This indicates that the method performs well. A key advantage in using the method is the speed with which designs are obtainable. This is shown in the Figure I attached. It can be seen that the time taken to obtain the optimal solution increases drastically as the network size increases (a characteristic of all NP-hard problems), and attempting to find the optimal solution for large network design problems is impractical.

Figure 1 illustrates the time taken by the proposed method to obtain a network topology design versus the network size. The time taken to obtain the optimal solution is also shown.

We claim:

1. A fast method for obtaining a low-cost, and possibly least cost, backbone network forming the basic infrastructure of an Internet Service Provider (ISP)

(a) consisting of the routers and the links administered and operated by the ISP – over a given set of router locations, given a traffic demand matrix, the said traffic originating from the demands of the ISP's customers, and a link distance – bandwidth cost matrix, the said cost matrix being given by a bandwidth provider, comprising

(b) providing sequential addition of traffic demands, each demand being routed over a minimum weight path, where all possible links given the router locations are considered and the link weights are assigned, for each demand by a differential cost analysis,

(c) ensuring that a demand between any ingress router and egress router (referred to as an "ingress-egress pair") is not split,

(d) ensuring symmetric routing of demands between any ingress-egress pair,

(e) ensuring hop-constrained routing, viz., that a demand between any ingress-egress pair traverses at most a pre-specified number of hops.

2. The method as claimed in claim 1, wherein the said method consists of two methods, named Link Removal Heuristic (LRH) and Demand Removal

Heuristic (DRH), that reduce cost by further refinements to the network, the said refinements being obtained by the following steps:

a) removal of a link (demand) from the network

b) updation of the network such that the removal link (demand) can be re-inserted into the network at least incremental cost.

3. The method as claimed in claim 1 wherein the said method is started by Surviving Single Link Failures Heuristic (SSLFH)method which incorporates survivability constraints in the design and thereby ensures sufficient spare capacity that is sued in the event of single link failures, the said method consisting of a repeated application of the following steps:

a. removal of a link from the network,

b. re-routing of all demands that were routed through the above link such that the re-routing can be accomplished at least incremental cost.

5. The methods as claimed in claims 1, 2 & 3 further including

a) the placement of routers

b) alternate paths for connections passing through a failed link.

Dated this day of March 2003.

AGARWAL JETLEY & CO.
A-2/78, SAFDARJUNG ENCLAVE,
NEW DELHI 110029
AGENTS FOR THE APPLICANTS

To
The Controller of Patents,
The Patents Office, New Delhi

13

12 MAR 2003

ABSTRACT

Network design and management are two fundamental problems faced by Internet Service Providers (ISPs). The objective of the invention is to obtain a network that costs the least while supporting the predicted traffic requirements. Given n node, there can be at most n(n-1) directed links in the network. Each potential link has two variables associated with it: current_capacity and current_traffic. The method is as follows:

Obtain an ordered sequence of (source, destination) pairs. For each potential link, consider the minimum capacity that would be required if the predicted traffic for that pair was made to flow through that link. For every link on the selected path, update the current_capacity and the current_cost. The aim of the method is to augment current topology such that the predicted traffic requirement between the chosen (source destination) pair can be stopped with minimum additional cost. There are two methods to further improve the cost of topology. The first method is "Link Removal Heuristic" (LRH) and the second is "Flow Removal Heuristic" (FRH). A key advantage of using the invented method is the speed with which designs are obtainable. The time taken to obtain the optimal solution increases drastically as the network size increases and attempting to find the optimal solution for large design problems is impractical.

ORIGINAL

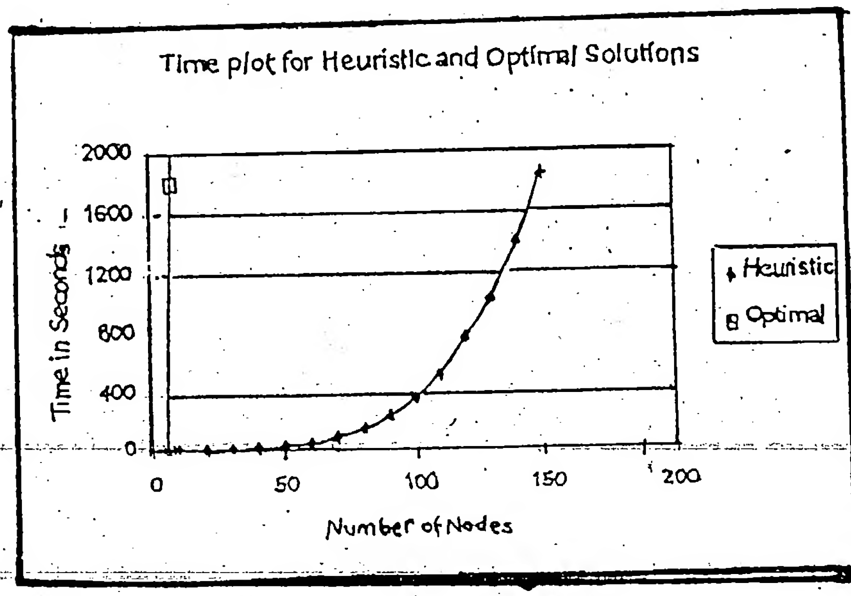

FIG.-1

FOR THE APPLICANT

PRAVEEN AGGARWAL